arXiv:1401.3518v1 [stat.AP] 15 Jan 2014

Percolation under Noise: Detecting Explosive Percolation Using the Second Largest Component

Wes Viles*,[1] Cedric E. Ginestet*,[1,*] Ariana Tang,[1] Mark A. Kramer,[1] and Eric D. Kolaczyk[1]

[1] *Department of Mathematics and Statistics, Boston University*[†]

(Dated: January 16, 2014)

We consider the problem of distinguishing classical (Erdős-Rényi) percolation from explosive (Achlioptas) percolation, under noise. A statistical model of percolation is constructed allowing for the birth and death of edges as well as the presence of noise in the observations. This graph-valued stochastic process is composed of a latent and an observed non-stationary process, where the observed graph process is corrupted by Type I and Type II errors. This produces a hidden Markov graph model. We show that for certain choices of parameters controlling the noise, the classical (ER) percolation is visually indistinguishable from the explosive (Achlioptas) percolation model. In this setting, we compare two different criteria for discriminating between these two percolation models, based on a quantile difference (QD) of the first component's size and on the maximal size of the second largest component. We show through data simulations that this second criterion outperforms the QD of the first component's size, in terms of discriminatory power. The maximal size of the second component therefore provides a useful statistic for distinguishing between the ER and Achlioptas models of percolation, under physically motivated conditions for the birth and death of edges, and under noise. The potential application of the proposed criteria for percolation detection in clinical neuroscience is also discussed.

I. INTRODUCTION

Understanding the emergence of organized structure in dynamic networks remains an active research area. In the study of random networks, percolation –the sudden emergence of a giant connected component (GCC)– is of critical importance from a theoretical, applied and statistical perspective. Percolation in the Erdős-Rényi (ER) model constitutes one of the first examples of a fully characterized mathematical phase transition [2, 7]. While the ER model of percolation is an example of a (second order) continuous phase transition, recent efforts have focused on identifying the conditions under which a random network process can yield a (first order) discontinuous percolation [18].

One of most popular attempts to model discontinuous percolation has been the Achlioptas' process and its variants [1]. The Achlioptas' product rule (PR) slows down the growth of the GCC by favoring the creation of edges between small connected components. Although this particular percolation model has been shown to be, in fact, continuous and therefore of second order [17, 18]; it nonetheless provides an interesting alternative to the ER model. Achlioptas' processes have indeed generated a substantial amount of theoretical work, whereby authors have explored related strategies for producing explosive percolation in random networks [3, 4, 14]. In addition, Riordan and Warnke [18] have shown that genuine first-order phase transitions can be realized by systematically adding, at every step of the process, the edge that joins the two smallest components in the *entire* network.

Interest in network percolation has been fueled by its relevance to several application domains. In clinical neuroscience, for instance, epileptic seizures have been associated with the sudden emergence of coupled activity across the brain [10, 11, 16, 20–22]. The resulting functional networks –in which edges indicate strong enough coupling between brain regions [19]– are consistent with the notion of percolation. A better understanding of the type of phase transitions undergone at different stages of the seizure, may aid in the development of novel strategies for the treatment of epilepsy [13].

The rich theory on percolation, and its application to real world data, motivates the following question: How can we distinguish between different percolation regimes in practice? Previous theoretical work has concentrated on noise-free percolation, which constitutes an idealized perspective on percolation processes. In practice, however, the sampling of real-world networks is likely to be corrupted by measurement errors. Moreover, network growth has generally been conceived as a monotonic process, whereby only edge creations are allowed. However, this assumption may be too restrictive, since in real-world networks, the number of edges may increase and decrease over time, in a stochastic manner (see example in Figure 1). Finally, to the best of our knowledge, there does not currently exist a statistical framework for distinguishing between different types of percolation regimes in the presence of edge birth and death, and noise.

In this paper, we propose a framework to distinguish between different percolation regimes in practice. To do so, we formulate the problem of recognizing a per-

---

* The authors marked ∗ have equally contributed to the production of this article. CEG is the corresponding author.

† EDK and CEG were supported by a grant from the Air Force Office for Scientific Research (AFOSR), whose grant number is FA9550-12-1-0102. MAK was supported by a Career Award at the Scientific Interface from the Burroughs Wellcome Fund. The authors thank Dr. Sydney Cash for providing the brain voltage data analyzed in Figure 1.

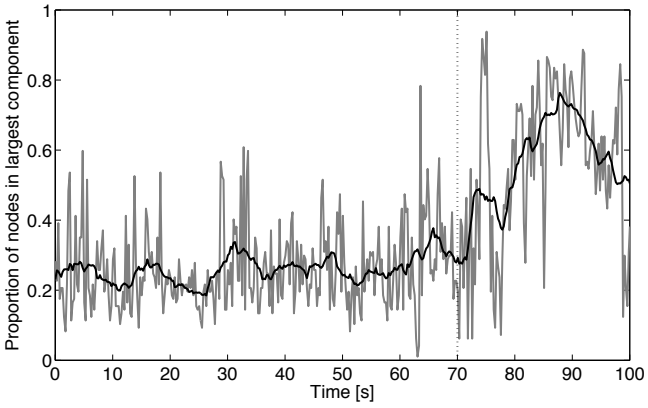

FIG. 1. Proportion of nodes in the largest component as a function of time for a functional network deduced from the electrocorticogram of a single patient with epilepsy during a seizure [see 12, for details]. The weighted functional networks have been binarized by conducting independent hypothesis tests on the maximum absolute values of the cross-correlations over $0.5s$ windows with 50% overlap and after correcting for multiple comparisons. The black trace is a smoothed version of this process. Seizure onset was clinically estimated to occur at the vertical dotted line.

colation regime from noisy observations as a question of statistical inference. Under this framework, we compare the discriminatory power of two potential percolation features deduced from the evolution of the first and second component of an observed dynamic network. We test this framework in simulation by constructing a hidden Markov graph model, which encompasses both a non-stationary latent process characterized by birth and death of edges, and an observed graph process that introduces both Type I and Type II errors. We show that edge death and noise make a statistic deduced from the first component ineffective in distinguishing between the standard ER second-order percolation and Achlioptas' explosive percolation. However, a different detection criterion –based on the size of the second component– successfully discriminates between the two percolation regimes in the presence of edge death and noise. These results provide a framework for distinguishing percolation regimes in practice.

II. PERCOLATION MODELS

A. Birth/Death Erdős-Rényi (ER) Process

We first construct a graph-valued stochastic process that exhibits the Markov property. This provides a realistic model for generating noisy percolation processes, while maintaining a sufficient level of computational tractability. We will denote a sequence of graph-valued random variables on n vertices by

$$\big\{G_t = (V, E_t) : t = 0, \ldots, T\big\}. \tag{1}$$

At each time step, a single edge is either added or deleted. Such a sequence will be said to be *Markov* if its edge sets are controlled by a Markov chain. We impose this dependence through the use of a binary random variable, denoted $\{Y_t : t = 0, \ldots, T\}$, whose state space is $\{0, 1\}$. This Markov chain is characterized by the following transition probability matrix P, for some choices of the birth and death rates, denoted respectively by p and q, and taking values in $[0, 1]$.

	$Y_{t+1} = 0$	$Y_{t+1} = 1$
$Y_t = 0$	$1 - p$	p
$Y_t = 1$	q	$1 - q$

Following customary notation, the entries of P will be denoted by $\mathbb{P}[Y_{t+1} = j | Y_t = i]$, with rows summing to one. The graph-valued Markov chain, G_t, is then obtained by associating Y_t with the addition or deletion of an edge in each edge set, E_t. Thus, provided that $p, q \neq 0$, it follows that the state space of this graph-valued Markov chain is the space of all simple graphs on n vertices, since every graph is reachable with positive probability.

In general, p and q are not required to sum to one. It will be of interest to let $p > q$ in order to study the large-scale behavior of the G_t's as the graph process accumulates edges. Moreover, observe that Y_t is a (time) *homogeneous* Markov chain, since $\mathbb{P}[Y_{t+1} = \omega | Y_t = \omega'] = \mathbb{P}[Y_1 = \omega | Y_0 = \omega']$, for any $\omega, \omega' \in \{0, 1\}$ and every t.

Now, suppose that there exist $m_t := |E_t|$ edges at time t in G_t and let $X_t(e)$ denotes the 'status' of edge e at time t, such that $X_t(e) = 1$, if that edge is present and $X_t(e) = 0$, otherwise. Note that we have here two different sources of dependence. On one hand, the edges are dependent on each other, since no more than one edge can be added or deleted at every time step. On the other hand, the edges are also dependent over time, since the status of an edge at time $t + 1$ depends on the status of that same edge at time t.

In the sequel, we will concentrate on a special case of this birth and death process, where we will set $p = 1 - q$. This leads to simplified marginal distributions for the edges. Additional details of this birth and death model are provided in appendix A.

B. Birth/Death Product Rule (PR) Process

We extend the standard Achlioptas' framework of PR percolation to a birth and death process, by devising death steps. This model is analogous to the aforementioned ER birth/death model, except for the choice of the probability distribution of the latent $X_t(e)$'s. As for the ER birth/death model, a binary random variable, Y_t, controls the addition or deletion of edges in each G_t. However, in the case of the PR model, the choice of the edge to be added or to be deleted is not uniform over the E_t's. Here, this choice depends on the modular structure of the graph at time t. Therefore, as for the ER model, we obtain a non-stationary stochastic process.

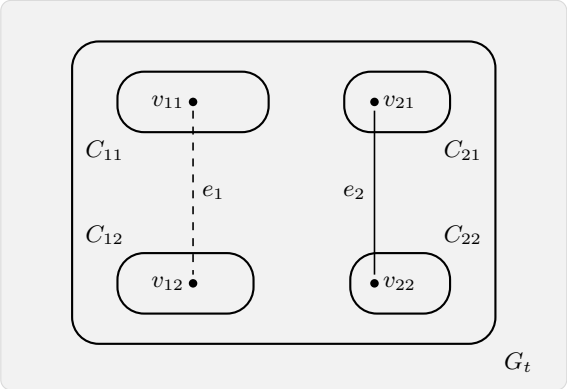

FIG. 2. Birth/death steps for the product rule (PR) in the Achlioptas' model of percolation. In this example, edge e_2 was born *before* edge e_1, since $|C_{11}||C_{12}| > |C_{21}||C_{22}|$. Therefore, when e_1 and e_2 are selected during a death step, e_2 is discarded *after* e_1. This death step specification ensures that births and deaths constitute genuine reverse PR operations.

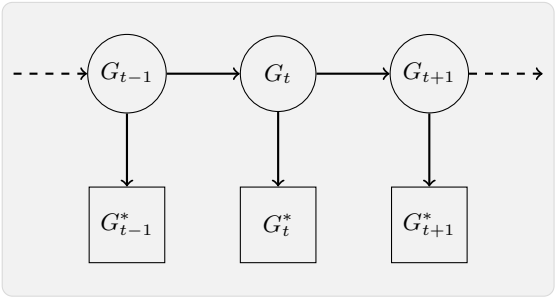

FIG. 3. Directed Acyclic Graph (DAG) representation of the hidden Markov process combining a *latent* stochastic graph process in the first row denoted by G_t, with an *observed* stochastic graph process contaminated by noise in the second row, denoted by G^*_t. Directed arrows indicate probabilistic dependence, such that the distribution of the observed G^*_t depends on the value taken by the latent graph, G_t.

Assuming that $Y_t = 1$, the addition of a new edge is conducted by uniformly choosing two candidate vertex pairs among all the edges in E_t^C, the complement of the edge set, E_t. These two candidate edge pairs are denoted by $e_1 := (v_{11}, v_{12})$ and $e_2 := (v_{21}, v_{22})$, and satisfy $X_t(e_1) = 0$ and $X_t(e_2) = 0$, since $e_1, e_2 \in E_t^C$, as in Figure 2. We then evaluate the size of the connected components to which v_{11}, v_{12}, v_{21} and v_{22} belong. These four connected components are denoted by C_{11}, C_{12}, C_{21}, and C_{22}, respectively. Then, following Achlioptas et al. [1], we apply the following product rule: If $|C_{11}||C_{12}| < |C_{21}||C_{22}|$, then $X_{t+1}(e_1) = 1$; otherwise, $X_{t+1}(e_2) = 1$.

Conversely, the death or deletion of an edge is handled in a symmetric manner. When $Y_t = 0$, we uniformly select two candidate edges from E_t. These vertex pairs are denoted $e_1 := (v_{11}, v_{12})$ and $e_2 := (v_{21}, v_{22})$ and satisfy $X_t(e_1) = 1$ and $X_t(e_2) = 1$, since $e_1, e_2 \in E_t$. Next, we set $X_t(e_1) = 0$ and $X_t(e_2) = 0$, in order to compute the size of the connected components to which v_{11}, v_{12}, v_{21} and v_{22} would belong to, if these edges were absent. This is done in order to ensure that the deletion of an edge exactly corresponds to the reverse operation of the addition of an edge under PR. Now, after having deleted these edges and computed the sizes of C_{11}, C_{12}, C_{21}, and C_{22}; we decide which edge should re-enter G_t, in order to produce G_{t+1}. Such a decision is also based on the PR, such that if $|C_{11}||C_{12}| < |C_{21}||C_{22}|$, then $X_{t+1}(e_2) = 0$; otherwise, $X_{t+1}(e_1) = 0$.

This choice of specification for the death step ensures that the *ordering* of the creation and deletion of edges are symmetrical. Given a sequence of two edges $\{e_1, e_2\}$ successively born during two time steps of G_t, if we encounter a death step, where both e_1 and e_2 are selected, we would then delete these edges in the reverse order, by eliminating e_2 before e_1. This order-preserving property is illustrated in Figure 2. This constraint ensures that births and deaths are genuine reverse PR operations. In addition, observe that, as for the ER percolation process, this chain is *irreducible*, in the sense that there is positive probability of transitioning from any given edge configuration to any other in the space of the edge sets of G.

C. Hidden Markov Graph Model

Next, we assume that there exists a time-independent error process, which produces at each time point an *observed* edge status $X^*_t(e)$. This stochastic process is governed by two additional parameters α and β, whose behavior can be described using a traditional 'confusion matrix', such that for any $\alpha, \beta \in [0, 1]$, we have

	$X^*_t(e) = 0$	$X^*_t(e) = 1$
$X_t(e) = 0$	$1 - \alpha$	α
$X_t(e) = 1$	β	$1 - \beta$

The $X_t(e)$'s and $X^*_t(e)$'s are here treated as *latent* and *observed* stochastic processes, respectively, and α and β can therefore be interpreted as the Type I (false positive) and Type II (false negative) error probabilities. Combining the graph-valued Markov latent process with this time-independent error process, we obtain a graph-valued *hidden Markov process*, as described in figure 3. From this schematic representation, one can immediately see that the observed graphs denoted G^*_{t-1}, G^*_t and G^*_{t+1} are conditionally independent, given the latent graph process, G_t.

For the ER model, under the assumption that $p = 1-q$, these two stochastic graph processes can be combined by taking into account the time-dependence of the $X^*_t(e)$'s. In this case, the corresponding transition matrix linking the observed and latent processes is available in closed-form. Details of these derivations are provided in appendix A.

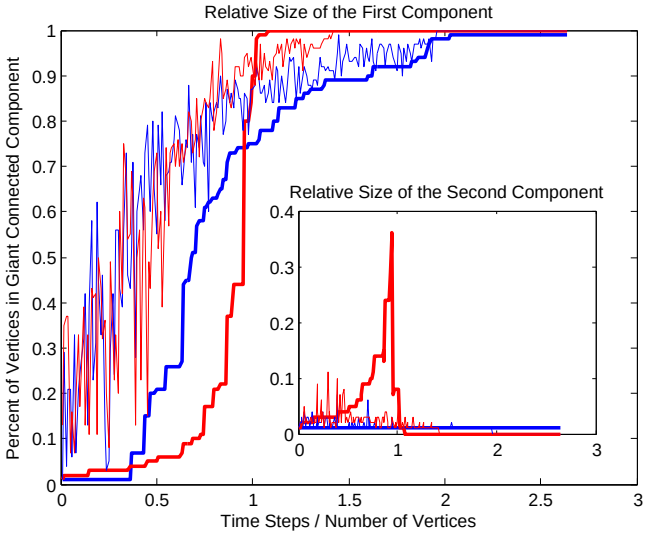
Relative Size of the First Component
Percent of Vertices in Giant Connected Component
Time Steps / Number of Vertices
Relative Size of the Second Component

FIG. 4. Percentage of vertices in the giant connected component (GCC) in the main window, and the corresponding sizes of the second largest component in inset, for a birth rate of $p = 1$ and death rate of $q = 0$, for the **ER (blue)** and **PR (red)** percolation models. In bold, these results are reported for a noise-free model, whereas the thin lines represent noisy simulations with Type I/II error rates of $\alpha = .0125$, $\beta = .01$. Observe that the two models become almost indistinguishable when noise is added to these processes.

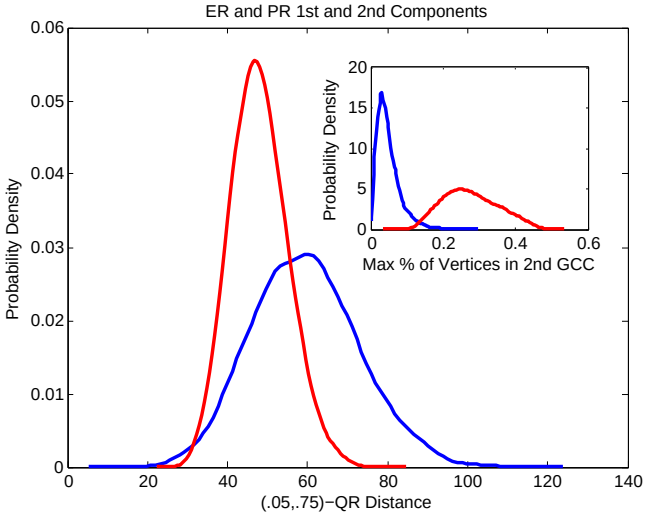

FIG. 5. Densities of the QD(.05, .75) of the first component's size in the main window, and densities of the maximal size of the second component (inset) are reported for the **ER (blue)** and **PR (red)** percolation models, based on 1,000 draws from the distributions of these two (noise-free) models. Note that the difference in scales of the density values of the y-axes in these two figures is due to the difference in scales of the x-axes.

III. DETECTING EXPLOSIVE PERCOLATION

Explosive percolation is expected to produce a sharper phase transition than a typical ER percolation. When considering noisy observations, however, detecting such differences through visual inspection only is hard. Figure 4 illustrates this problem, by comparing noisy and noise-free graph sequences for both explosive and ER models. Beyond visual inspection, the problem of discriminating between these two models of percolation can be formulated as a hypothesis-testing problem: The null hypothesis, denoted H_0, states that the observed process corresponds to an ER percolation, whereas the alternative hypothesis, H_1, is that the observed process does not correspond to this type of percolation model.

To proceed with this hypothesis-testing problem in practice, we specify a population parameter summarizing the percolation process, say θ^{ER} and θ, for the ER and target models, respectively. This leads to a hypothesis test of the form,

$$H_0 : \theta^{\text{ER}} = \theta, \qquad \text{and} \qquad H_1 : \theta^{\text{ER}} \neq \theta.$$

Several population parameters could be used for the purpose of discriminating between these two models of percolation. A natural candidate for such parameters would be a measure of the sharpness of the transition of the first component's size. The main panel of Figure 4, however, suggests that this population parameter will not have sufficient discriminatory power, when confronted with a substantial amount of observational noise.

Therefore, as a second candidate population parameter, we consider the following natural extension: The size of the *second largest* component. This appears to provide a more sensitive marker of the sharp phase transition exhibited by explosive percolation models. This 'second-order' property is motivated by the mechanism of the product rule, which underlies the Achlioptas' graph process, which we will use as representative of the alternative hypothesis, H_0 (i.e., reducing our testing problem to the comparison of two point null hypotheses). In what follows, we will show that the use of the size of the second largest component as a statistical marker to distinguish the two percolation regimes exhibits greater discriminatory power, than a statistic solely based on the first component.

The differences between the candidate percolation models are therefore quantified using two criteria: (i) a quantile difference (QD) of the distribution of the size of the GCC, and (ii) the maximal size of the second component over the entire time period. These two criteria are formally defined as follows. Given the graph process, $G_t = (V, E_t)$, and denoting the vertex subset of the largest component in G_t by $S_{1,t}$, we define the cumulative edge function as the *cardinality* of $S_{1,t}$, normalized by the maximal number of edges in the graph, such that

$$F(t) := \frac{|S_{1,t}|}{\binom{n}{2}}.$$

Although this function is not a cumulative distribution function (CDF), one can nonetheless uniquely define

quantiles using the standard definition of quantiles for the CDFs of discrete random variables; such that for any $x \in [0, 1]$, we have

$$Q(x) := \min_{t=1,\dots,T} \{t : F(t) \geq x\},$$

where T is the maximal number of time steps in the graph process. In this paper, we are especially interested in quantile differences of the following form, which can be treated as a generalization of the classical interquartile range,

$$\mathrm{QD}(x_1, x_2) := Q(x_2) - Q(x_1).$$

One can observe from Figure 1 that the size of the GCC, in the functional networks associated with a seizure, rarely exceeds 80%. Therefore, we have here adopted a quantile difference that reflects the range of the distribution of the GCC in this practical setting. Thus, the criterion of interest will be the following, $\theta_{\mathrm{QD}} := Q(.75) - Q(.05)$. This parameter quantifies the steepness of the phase transition, the larger the QD criterion, the longer the transition to a fully connected graph.

As a second criteria to distinguish the two percolation regimes, we consider the maximal size attained by the second largest component over the entire time period of the dynamic network observation. If one defines the vertex set of the second largest component at time t by $S_{2,t}$, this second criterion can be expressed as $\theta_{\mathrm{Sec}} = \max\{|S_{2,t}| : t = 1, \dots, T\}$. This quantity is expected to constitute a good marker of the steepness of the phase transition, since it reflects the extent of separation of the graph process into large connected subgraphs. Indeed, a direct consequence of the Achlioptas' construction rule is that by inhibiting the growth of a single large component, we necessarily increase the production of several subcomponents.

Statistical inference on these two criteria is then drawn using a Monte Carlo hypothesis test. Letting the parameter $\theta := \theta_{\mathrm{QD}}$, and selecting the candidate percolation to be drawn from a PR process, we consider the following null and alternative hypotheses,

$$H_0 : \theta^{\mathrm{ER}} = \theta^{\mathrm{PR}}, \qquad \text{and} \qquad H_1 : \theta^{\mathrm{ER}} > \theta^{\mathrm{PR}},$$

respectively. The direction of this test is justified by the fact that we expect explosive percolation to occur rapidly, and thus to exhibit a smaller amount of variability in the size of its GCC, when transitioning to a fully connected graph. Our second criterion, by contrast, is tested in the opposite direction, since we naturally anticipate the PR process to be characterized by a *larger* maximal second component. Thus, for $\theta := \theta_{\mathrm{Sec}}$, the alternative hypothesis becomes $H_1 : \theta^{\mathrm{ER}} < \theta^{\mathrm{PR}}$.

In the results reported in this paper, the distributions of the ER and PR graph processes are known. It therefore suffices to simulate from these densities in order to construct the distribution of the two test statistics at hand. This procedure is illustrated in figure 5. We are

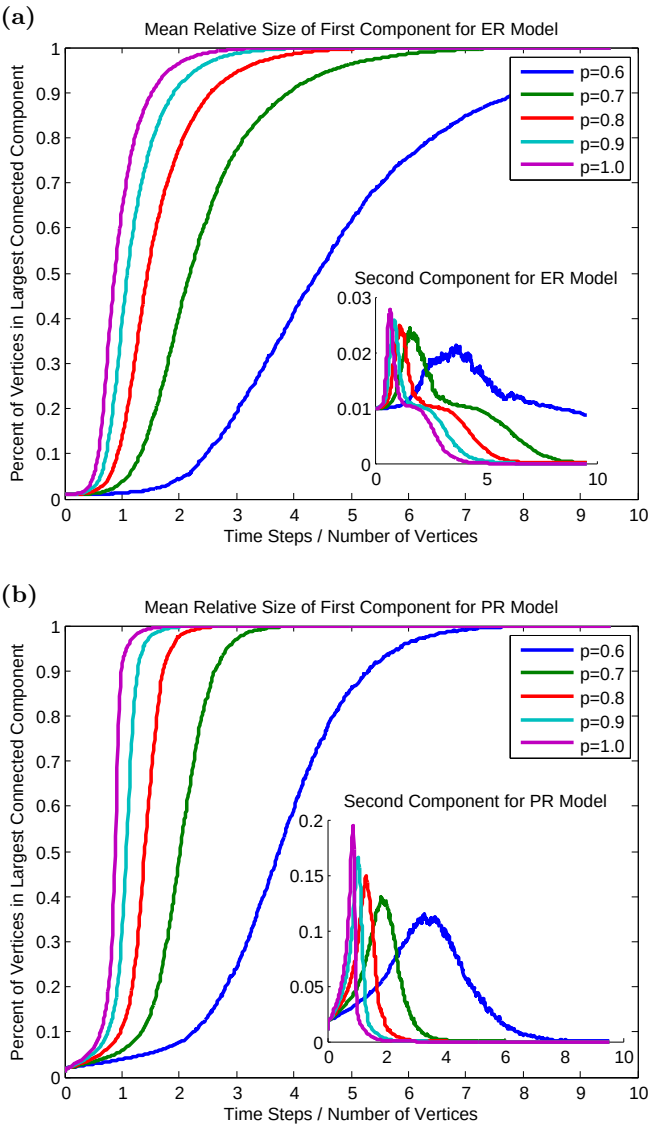

FIG. 6. Percentage of vertices in giant connected component (GCC) with respect to time, under the ER and PR models, in panels (a) and (b), respectively; for different choices of birth rates, p. In each panel, the corresponding sizes of the second largest component are also reported in inset. The edge death rate is here dependent on the birth rate, such that $q = 1 - p$. Each curve represents the mean of 1,000 different simulations.

especially interested in the discriminatory powers of these statistics, and we will therefore compare their respective merits, using the true positive and false positive rates, within a receiver operating characteristic (ROC) framework. For presentational convenience, the distributions of interest were smoothed using a normal density kernel, before computing the ROC curves and corresponding *areas under the curves* (AUCs). The computation of the AUCs allows us to summarize the differences between the models over the entire time period.

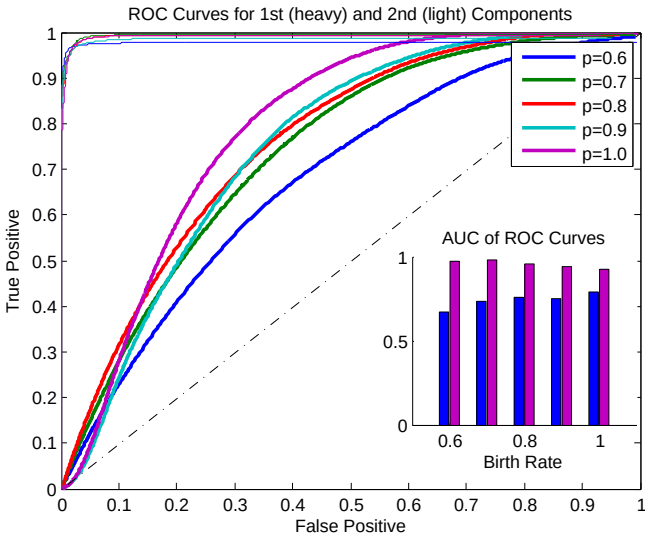

FIG. 7. Receiver operating characteristic (ROC) curves based on the QD of the first component's size (bold lines) and on the maximal size of the second component (thin lines). Each curve was produced using 1,000 simulations from each percolation model. In inset, the areas under the curves (AUCs) of these different ROC curves are compared for different choices of the birth rate, p, with the blue and purple bars denoting the ER and PR models, respectively.

IV. RESULTS

A. Birth/Death Processes

The ER and PR percolation models were simulated on graphs of $n = 100$ vertices. We first explored the effect of varying the birth and death rates on the behavior of the two statistical criteria of interest, under the ER and PR models. The results of these Monte Carlo simulations are reported in Figure 6, where each curve is the mean of 1,000 different synthetic data sets. In these simulations, the death rate was set to $q = 1 - p$, and therefore the value of p controls both the birth and death rates.

The main effect of a change in p is to delay percolation, and to diminish the steepness of the phase transition. Observe that as p decreases, percolation tends to occur at a later time step in both the ER and PR models (Figure 6). In particular, for the lowest birth rate that we investigated ($p = 0.6$), the ER model did not produce a fully connected graph within the number of iterations considered, as can be seen from Figure 6(a). When p was set to values equal to or less than 0.5, no phase transition could be observed, and these results are not reported.

The size of the second component was similarly affected by changes in p. Decreasing the birth rate delayed the time at which the size of the second component attained its highest value. Moreover, lower values of p also yielded second components with smaller maximal sizes, under both the ER and PR models. Interestingly, we note that the time points at which the second components reach a maximal size tend to coincide in both models. Thus, it would be difficult to distinguish between these two percolation models on the sole basis of the *timing* of the occurrence of the maximal size of the second components. By contrast, the relative maximal *size* of the second components in the ER and PR models differ by approximately one order of magnitude, thereby providing a natural criterion for discriminating between these two types of percolation, as can be observed by comparing the inset figures in 6(a) and 6(b).

We formally quantified these differences in discriminatory powers by studying the ROC curves of these two criteria under different choices of p (Figure 7). The maximal size of the second component substantially outperforms the relative size of the first component, for all values of p. The stark difference between these discriminatory criteria can be understood by considering the amount of overlap of the distributions of these two criteria in Figure 5. Whereas the distributions of the QD of the size of the GCC under the two models exhibit a large amount of overlap; the distributions of the maximal size of the second component, by contrast, share very little common support. These differences in support account for the substantive gains in discriminatory power by the maximal size of the second component, reported in Figure 7.

In addition, we note the QD of the size of the first component was more sensitive to choices of p than the maximal size of the second component. As p diminishes, it becomes increasingly more difficult to discriminate between the ER and PR percolation models, using the QD of the size of the first component. This suggests that this criterion is more sensitive to a non-zero death rate, than the maximal size of the second component, which provides further support for the use of this latter criterion, in practice.

B. Percolation under Noise

Secondly, we considered the effect of introducing noise in these models. The results reported in Figure 8 were produced using our proposed hidden Markov graph model, and are averaged over 1,000 simulations. We were especially interested in the effect of Type I and Type II errors on our ability to discriminate between classical and explosive percolation, using the two criteria under scrutiny. Both the Type I and Type II error rates were made to vary between 0 and .01.

From Figure 8, one can observe that the two types of errors had markedly different effects on the AUCs of the two discriminatory criteria. Introducing Type I errors led to a substantial diminution of the AUCs for both the QD of the size of the first component in (a), and the maximal size of the second component in (b). In particular, note that the two criteria reached equivalent levels of discriminatory power for $\alpha = 0.01$. Thus, although the maximal size of the second component remains a more

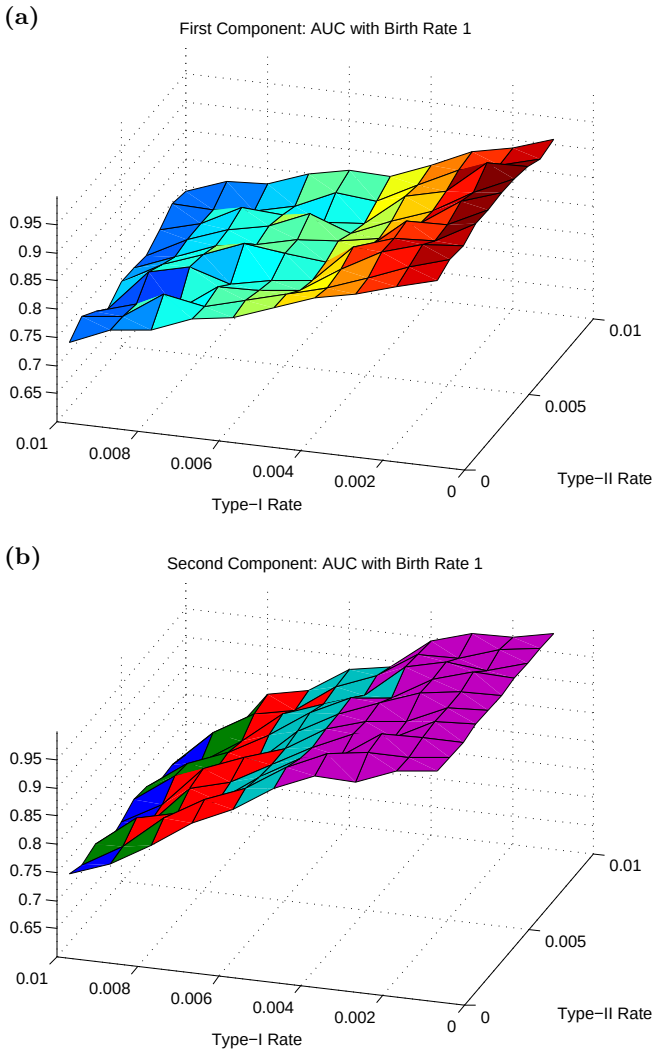

FIG. 8. Areas under the curve (AUCs) for ROC curves for 10^5 model simulations based on the QD of the first component's size in panel (a), and based on the maximal size of the second component in panel (b). We have varied the Type I and Type II error rates within a hidden Markov graph model framework. The birth/death rates have been kept constant with $p = 1$.

useful criterion for distinguishing between the ER and PR models than the QD of the size of the first component, these two criteria exhibit comparable performance, under a moderate amount of Type I error.

The impact of increasing the Type II error rate on the behavior of these two criteria was negligible. Introducing false negatives in the ER and PR models slightly increased the AUCs of both the QD of the first component's size, and the maximal size of the second component. Thus, large Type II error rates may be marginally advantageous for discriminating between these two models of percolation, under the scenarios studied.

V. DISCUSSION AND CONCLUSIONS

In this paper, we have extended existing models of percolation, by allowing for edge deletion steps and noisy observations. These modeling extensions have been articulated within a hidden Markov graph process, which builds links with the existing literature on the statistical properties of this family of models [5, 8, 15]. Moreover, we have compared different summary statistics for distinguishing between the ER and PR percolation models. Overall, for different birth and death rates, and for a range of noise levels, the maximal size of the second component was found to have greater discriminatory power than the QD of the size of the GCC.

Several methodological challenges remain before such models can be directly used for percolation detection on real-world data. Throughout this paper, we have considered the QD of the size of the first component, using a particular choice of quantiles for this discriminatory criterion. In practice, an optimal choice of quantiles for quantifying the steepness of such phase transitions may be motivated by different factors, including (i) the range of the observations, and (ii) the need for early detection. We discuss these two practical aspects, in turn.

Firstly, note that when considering real-world applications, we rarely observe fully connected networks. In the data reported in Figure 1, for instance, the size of the GCC encompasses at most 90% of the edges in the saturated network. The choice of the quantile interval of interest for the first component will be therefore automatically constrained by the range of the observations in the data at hand. Therefore, as in sequential detection analysis, the statistical objective is to detect the outcome, on the basis of as little data as possible. Such constraints would naturally lead to a relatively narrow quantile range.

Secondly, in the context of clinical neuroscience and with particular emphasis on the prevention of a seizure; the detection of a percolation regime may be linked with patients' health and survival. In such cases, early detection will usually be favored, as this is likely to be associated with desirable clinical outcomes. Explosive percolation, such as the Achlioptas' PR process studied in this paper, is consistent with the sudden manifestation of a seizure as a highly synchronized event. Classifying models of percolation may then be utilized to deepen understanding of seizures in epilepsy, and a statistical identification of explosive phase transitions may facilitate immediate, targeted intervention (e.g., an electrical stimulus).

Further work in this area could be focused on estimating a percolation model from a given sequence of observed networks. In this sense, this work also contributes to the growing literature on time-indexed graph processes [9]. In such cases, the birth and death rates will need to be estimated, as well as the Type I and Type II error probabilities. These different parameters may not be fully identifiable from the data, and further constraints are likely

to be necessary, in order to discriminate between the two percolation models considered in this paper. Such estimation, however, may be amenable to a Bayesian formulation, as commonly implemented for hidden Markov models [6].

Appendix A: Details of Birth/Death ER Process

Here, we describe the closed-form formulas of the probabilities of edge inclusion and edge deletion in the observed graph processes under the ER model. These analytic results are obtained by assuming that the birth and death probabilities are straightforwardly related, such that $p = 1 - q$. Such derivations may be useful for other authors, who may want to replicate these results, or extend the applications of the noisy model of percolation.

In this birth and death graph process, each edge is treated separately by integrating out the dependence of all other edges in the graph, and considering the marginal distribution of every $X_t(e)$. As before, we will here refer to $X_t(e)$ as the *latent* edge status, and $m_t := |E_t|$ will indicate the number of edges in the graph at time t. Given the Markov random variable Y_t, the conditional transition matrix for every $X_{t+1}(e)$, given some value of Y_t takes the following form,

	$X_{t+1}(e) = 0$	$X_{t+1}(e) = 1$
$X_t(e) = 0$	$\frac{\binom{n}{2} - m_t - I\{Y_t = 1\}}{\binom{n}{2} - m_t}$	$\frac{I\{Y_t = 1\}}{\binom{n}{2} - m_t}$
$X_t(e) = 1$	$\frac{I\{Y_t = 0\}}{m_t}$	$\frac{m_t - I\{Y_t = 0\}}{m_t}$

where $\binom{n}{2}$ denotes the number of edges in a saturated graph of size n, and where $I\{f(x)\}$ is the indicator function, which takes a value of 1 if $f(x)$ is true and 0, otherwise.

In this paper, we have concentrated on a special case of this birth/death process, where we have set $p = 1 - q$. This choice of p and q leads to the following characterization of the Y_t process,

$$\begin{aligned} \mathbb{P}\left[Y_{t+1} = 1\right] &= \mathbb{P}\left[Y_{t+1} = 1 | Y_t = 0\right] \\ &= \mathbb{P}\left[Y_{t+1} = 1 | Y_t = 1\right] = p, \end{aligned} \tag{A1}$$

and similarly, $\mathbb{P}\left[Y_{t+1} = 0\right] = 1 - p$. Under this simplifying assumption, the preceding conditional transition matrix becomes

	$X_{t+1}(e) = 0$	$X_{t+1}(e) = 1$
$X_t(e) = 0$	$\frac{\binom{n}{2} - m_t - p}{\binom{n}{2} - m_t}$	$\frac{p}{\binom{n}{2} - m_t}$
$X_t(e) = 1$	$\frac{1-p}{m_t}$	$\frac{m_t - 1 + p}{m_t}$

Each entry is obtained by taking the expectation with respect to Y_t. That is, $\mathbb{P}[X_{t+1}(e) = \omega | X_t(e) = \omega'] = \mathbb{E}[\mathbb{P}[X_{t+1}(e) = \omega | X_t(e) = \omega', Y_t]]$, for every $\omega, \omega' \in \{0, 1\}$, and where the marginal distribution of Y_t is known from equation (A1).

One can now combine the noise process described in section II C, with the birth and death stochastic process in order to link the latent and observed parts of the Markov hidden model. This gives the following table,

	$X^*_{t+1}(e) = 0$	$X^*_{t+1}(e) = 1$
$X_t(e) = 0$	$(1 - \alpha)\left(\frac{\binom{n}{2} - m_t - p}{\binom{n}{2} - m_t}\right)$	$\alpha\left(\frac{p}{\binom{n}{2} - m_t}\right)$
$X_t(e) = 1$	$\beta\left(\frac{1-p}{m_t}\right)$	$(1 - \beta)\left(\frac{m_t - 1 + p}{m_t}\right)$

Since this transition matrix links the latent and observed stochastic processes, one can immediately derive the marginal probabilities of the $X^*_t(e)$'s, such that

$$\mathbb{P}[X^*_{t+1}(e) = 0] = (1 - \alpha)\left(\frac{\binom{n}{2} - m_t - p}{\binom{n}{2} - m_t}\right) + \beta\left(\frac{1-p}{m_t}\right),$$

and similarly for $\mathbb{P}[X^*_t(e) = 1]$.

The resulting X_t process is *non-stationary*. Moreover, this also holds when considering the case $p = 1 - q$. Indeed, since the probability of adding a new edge at time $t + 1$ is dependent on the number of existing edges, $m_t := |E_t|$, at time t; it follows that the resulting joint distribution of any subset of the X_t's depends on the choice of t.

[1] Achlioptas, D., D'Souza, R.M., and Spencer, J. (2009). Explosive percolation in random networks. *Science*, **323**(5920), 1453–1455.

[2] Alon, N. and Spencer, J. (2004). *The Probabilistic Method*. Springer, London.

[3] Beveridge, A., Bohman, T., Frieze, A., and Pikhurko, O. (2007). Product rule wins a competitive game. *Proc. Amer. Math. Soc*, **135**, 3061–3071.

[4] Bohman, T. and Frieze, A. (2001). Avoiding a giant component. *Random Structures and Algorithms*, **19**, 75.

[5] Casteigts, A., Flocchini, P., Quattrociocchi, W., and Santoro, N. (2012). Time-varying graphs and dynamic networks. *International Journal of Parallel, Emergent and Distributed Systems*, **27**(5), 387–408.

[6] Ephraim, Y. and Merhav, N. (2002). Hidden markov processes. *IEEE Transactions on Information Theory*, **48**(6), 1518–1569.

[7] Erdős, P. and Rényi, A. (1960). The evolution of random graphs. *Publ. Math. Inst. Hungar. Acad. Sci.*, **5**, 17.

[8] Flocchini, P., Mans, B., and Santoro, N. (2012). On the exploration of time-varying networks. *Theoretical Computer Science*, **469**, 53–68.

[9] Grindrod, P. and Higham, D. (2009). Evolving graphs: Dynamical models, inverse problems and propagation.

Proceedings of the Royal Society, Series A, **0456**, 1–18.

[10] Guye, M., Regis, J., Tamura, M., Wendling, F., Gonigal, A.M., Chauvel, P., and Bartolomei, F. (2006). The role of corticothalamic coupling in human temporal lobe epilepsy. *Brain*, **129**(7), 1917–1928.

[11] Kramer, M., Eden, U., Kolaczyk, E., Zepeda, R., Eskandar, E., and Cash, S. (2010). Coalescence and fragmentation of cortical networks during focal seizures. *Journal of Neuroscience*, **30**, 10076–10085.

[12] Kramer, M., Eden, U., Cash, S., and Kolaczyk, E. (2009). Network inference with confidence from multivariate time series. *Phys. Rev. E.*, **79**, 061916.

[13] Kramer, M.A. and Cash, S.S. (2012). Epilepsy as a disorder of cortical network organization. *The Neuroscientist*, **18**(4), 360–372.

[14] Krivelevich, M., Lubetzky, E., and Sudakov, B. (2010). Hamiltonicity thresholds in Achlioptas processes. *Random Struct. Alg.*, **37**(1), 1–24.

[15] Lentz, H.H.K., Selhorst, T., and Sokolov, I.M. (2013). Unfolding accessibility provides a macroscopic approach to temporal networks. *Phys. Rev. Lett.*, **110**(11), 118701.

[16] Ponten, S., Bartolomei, F., and Stam, C. (2007). Small-world networks and epilepsy: graph theoretical analysis of intracerebrally recorded mesial temporal lobe seizures. *Clinical neurophysiology*, **118**(4), 918–927.

[17] Riordan, O. and Warnke, L. (2012). Achlioptas process phase transitions are continuous. *Annals of Applied Probability*, **4(22)**, 1450–1464.

[18] Riordan, O. and Warnke, L. (2011). Explosive percolation is continuous. *Science*, **333**(6040), 322–324.

[19] Rubinov, M. and Sporns, O. (2010). Complex network measures of brain connectivity: Uses and interpretations. *Neuroimage*, **52**, 1059–1069.

[20] Schindler, K., Amor, F., Gast, H., Muller, M., Stibal, A., Mariani, L., and Rummel, C. (2010). Peri-ictal correlation dynamics of high-frequency (80–200Hz) intracranial EEG. *Epilepsy research*, **89**(1), 72–81.

[21] Schindler, K., Elger, C.E., and Lehnertz, K. (2007). Increasing synchronization may promote seizure termination: evidence from status epilepticus. *Clinical neurophysiology*, **118**(9), 1955–1968.

[22] Schindler, K., Leung, H., Elger, C.E., and Lehnertz, K. (2007). Assessing seizure dynamics by analysing the correlation structure of multichannel intracranial EEG. *Brain*, **130**(1), 65–77.